Preprint typeset using LaTeX style emulateapj v. 5/2/11

arXiv:1403.2720v3 [astro-ph.CO] 9 Dec 2014

MEASURING THE POWER SPECTRUM OF DARK MATTER SUBSTRUCTURE USING STRONG GRAVITATIONAL LENSING

Yashar Hezaveh[1,2,3], Neal Dalal[4], Gilbert Holder[3], Theodore Kisner[5], Michael Kuhlen[6,7], Laurence Perreault Levasseur[8]

ABSTRACT

In recent years, it has become possible to detect individual dark matter subhalos near images of strongly lensed extended background galaxies. Typically, only the most massive subhalos in the strong lensing region may be detected this way. In this work, we show that strong lenses may also be used to constrain the much more numerous population of lower mass subhalos that are too small to be detected individually. In particular, we show that the power spectrum of projected density fluctuations in galaxy halos can be measured using strong gravitational lensing. We develop the mathematical framework of power spectrum estimation, and test our method on mock observations. We use our results to determine the types of observations required to measure the substructure power spectrum with high significance. We predict that deep observations (~ 10 hours on a single target) with current facilities can measure this power spectrum at the 3σ level, with no apparent degeneracy with unknown clumpiness in the background source structure or fluctuations from detector noise. Upcoming ALMA measurements of strong lenses are capable of placing strong constraints on the abundance of dark matter subhalos and the underlying particle nature of dark matter.

Subject headings: dark matter — gravitational lensing — galaxies: dwarf — galaxies: structure

1. INTRODUCTION

The abundance of substructure within the dark matter halos surrounding galaxies has been an area of intensive study for over a decade (e.g Moore et al. 1999; Klypin et al. 1999; Dalal & Kochanek 2002; Kravtsov et al. 2004; Simon & Geha 2007; Strigari et al. 2007). Dark matter substructure in the present-day universe is sensitive to the spectrum of primordial density fluctuations, generated in the very early universe, implying that a precise quantification of substructure can help constrain the physics of cosmic inflation (e.g., Viel et al. 2004). In addition, the microphysics of dark matter, such as its temperature or the strength of its interactions, can also influence the structure of dark matter halos and subhalos (e.g., Lovell et al. 2012; Rocha et al. 2013; Lovell et al. 2014). Studies of halo substructure can therefore probe multiple areas of fundamental physics. Comprehensive searches for faint and small satellite galaxies of the Milky Way have revealed that the number of low-mass observable satellites is significantly lower than what is predicted in CDM simulations (Kravtsov 2010), referred to as the "Missing Satellite Problem". One possible explanation for this discrepancy is that, perhaps, large numbers of dark matter subhalos exist but are not observed because they are devoid of baryons, rendering them effectively invisible. If so, then the paucity of dwarf satellite galaxies is a problem for galaxy formation models to address. Another possibility is that the predicted subhalos simply do not exist, pointing to new physics in the dark matter sector or inflation (e.g., Lovell et al. 2012; Rocha et al. 2013; Lovell et al. 2014).

Measuring the structure of dark matter on sub-galactic scales can therefore shed light on the nature of dark matter and star formation in dwarf halos. An unambiguous characterization of the structure of DM halos on sub-galactic scales requires a purely gravitational detection method. Gravitational lensing is an effective tool for mapping out mass, even completely dark mass. Strong lenses, which produce multiple images of distant sources, are sensitive to the presence of small-scale subhalos in lensing galaxies (Mao & Schneider 1998). Subhalos can induce perturbations in nearby images while leaving more distant (in angular separation) images unaffected. An analysis of the differences between multiple images of a background source can then reveal the presence of density perturbations near images. Dalal & Kochanek (2002) used the anomalous flux ratios of multiple images of 7 lensed radio quasars to constrain $f_{\rm sub}$, the fraction of galaxy mass locked in subhalos, finding $0.6\% < f_{\rm sub} < 7\%$ at 90% confidence. More recently, Vegetti et al. (2010, 2012) showed that galaxy-galaxy strong lenses can be used to detect subhalos, with two detections reported to date, resulting in $1.5\% < f_{\rm sub} < 6.9\%$ at 68% confidence for an assumed subhalo mass function $dn/dM \propto M^{-1.8}$ (Xu et al. 2013). In addition, Nierenberg et al. (2014) reported the detection of a subhalo in narrow-line emission of an optical quasar. These studies, however, have large uncertainties due to

[1] Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, 452 Lomita Mall, Stanford, CA 94305-4085, USA
[2] Department of Physics, Stanford University, 452 Lomita Mall, Stanford, CA 94305-4085, USA
[3] Department of Physics, McGill University, 3600 Rue University, Montreal, Quebec H3A 2T8, Canada
[4] Astronomy Department, University of Illinois at Urbana-Champaign, 1002 W. Green Street, Urbana IL 61801, USA
[5] Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
[6] Theoretical Astrophysics Center, University of California, Berkeley, CA 94720, USA
[7] LendUp, 237 Kearny St #372, San Francisco, CA 94108, USA
[8] DAMTP, University of Cambridge, Cambridge, CB3 0WA, United Kingdom

small sample size (e.g., lensed radio quasars) and/or limited sensitivity to subhalos (in case of extended source, e.g., galaxy-galaxy lenses). Improving this measurement requires *both* a significant increase in the sample of lenses, and a significant increase in the sensitivity of each lens to the effects of substructure. Hezaveh et al. (2013a) suggested that the spatial and spectral resolution of ALMA and the high signal to noise ratios of ALMA observations can allow us to detect of order one subhalo in bright lensed sub-mm lensed galaxies. The analysis of the large number of newly discovered systems in this sample (e.g. Hezaveh et al. 2013b; Vieira et al. 2013; Bussmann et al. 2013) has the potential to yield a high significance measurement of the abundance of subhalos with $M \gtrsim 10^8 M_\odot$.

As discussed in Hezaveh et al. (2013a), only the few most massive subhalos can be individually detected and characterized in typical strong lenses with extended sources. A much larger number of subhalos are expected to exist at lower masses (e.g. $\sim 10^6 M_\odot$), where various dark matter candidates give rise to drastically different subhalo abundances. Although these subhalos cannot be individually detected, they can collectively induce observable image perturbations. These collective perturbations allow the possibility of a *statistical* detection of the population of low mass subhalos. Dalal & Kochanek (2002) presented a method for statistically constraining the properties of the DM subhalo population using an ensemble of strong lensing systems. Their method, however, is computationally intensive, and therefore challenging to apply to the large data sets that will be provided for extended lensed images by instruments such as ALMA.

In this paper, we present an alternative method for constraining the population of undetected subhalos using strong lensing data. Instead of modeling individual subhalos, we show that it is possible to measure the power spectrum of the substructure density field by measuring the correlation of image residuals, after modeling the data with a lens with a smooth-potential. This technique allows us to probe the large population of unresolved subhalos with masses $\lesssim 10^7 M_\odot$, by revealing their abundance and their average density profiles. In §2, we discuss how the substructure power spectrum is related to the properties of subhalos. In §3, we present the mathematical framework underlying our method. Then in §4 we apply our methods to mock observations, and show how well the underlying substructure power spectrum may be recovered. Finally, we conclude in §5 with a discussion of the implications of this work.

In all calculations we have assumed a cosmology with $\Omega_\Lambda = 0.734$, $\Omega_m = 0.267$, and $h = 0.71$, and assumed that the lens was placed at $z_d = 0.5$.

2. THE PROJECTED SUBSTRUCTURE POWER SPECTRUM IN DARK MATTER HALOS

The distribution of substructure in ordinary galactic halos has been studied extensively in the literature (e.g. Diemand et al. 2007b, 2008; Stadel et al. 2009; Navarro et al. 2010). N-body simulations of ΛCDM cosmologies indicate that galactic dark matter halos contain subhalos whose abundance approximately follows a power-law distribution, $dN/dM \sim M^{-\alpha}$ with $\alpha \sim 1.9$, extending to the mass resolution limits of the simulations. Internally, subhalos appear to have density profiles consistent with tidally truncated NFW profiles. The spatial distribution of subhalos within their hosts may be somewhat less concentrated than the radial distribution of dark matter, due to tidal stripping and destruction at small radii (Diemand et al. 2004).

The density fluctuations associated with substructure may be considered as a random field, superimposed on the smoothly varying background density profile of the host. This field is not Gaussian. However, because the subhalo mass function rises so quickly towards low mass, much of this non-Gaussianity is generated by the few most massive halos. We expect to be able to detect these few massive subhalos individually, using direct lens modeling techniques (e.g. Hezaveh et al. 2013). Below the detection sensitivity of these direct modeling techniques (e.g. $\lesssim 5 \times 10^7 M_\odot$) the number of subhalos is very large, reducing the non-Gaussianity of the density field.

If we assume that the density field is Gaussian, we can fully characterize it by its power spectrum. A useful way to understand the power spectrum is to use the halo model (Cooray & Sheth 2002). This decomposes the power spectrum into its contributions from subhalos of different masses. The ingredients of the halo model are (1) the subhalo mass function, (2) the internal density profiles of the subhalos, and (3) the distribution of subhalo locations. Although few subhalos are physically at small radii, they will occasionally randomly project onto the strong lensing region ($r \sim 5 - 10$ kpc). This implies that the projected number densities of subhalos will (on average) be nearly constant as a function of projected r. In other words, over the small regions probed by strong lensing, we can assume that subhalos have a Poisson distribution with nearly constant projected number density. We can also neglect correlations among subhalo locations. In the language of the halo model, this corresponds to neglecting the 2-(sub)halo term. The reason it is safe to neglect subhalo-subhalo correlations is that subhalos reside within the tidal gravitational field of their host. Subhalos that are not gravitationally bound to each other will follow orbits whose relative orbital phases wrap by order-unity angles within a few dynamical times, i.e. a timescale that is short compared to the Hubble time (Chamberlain et al. 2014). Therefore, although subhalos have significant spatial correlations when they are accreted onto their hosts, those correlations should quickly decay due to tidal gravity. The exception to this argument is sub-substructure, i.e. sub-subhalos that are gravitationally bound within larger subhalos. In general, however, sub-substructure comprises a very small fraction of the total mass in substructure, and we therefore neglect it in our calculations (Diemand et al. 2007a).

We can therefore write down the substructure power spectrum as an integral over the subhalo mass function, weighted by the (Fourier transform of the) subhalo density profile, i.e. the 1-(sub)halo term:

$$P_\kappa(k) = \int \frac{dn}{dM} |\kappa_M(k)|^2 dM, \quad (1)$$

where $\kappa_M(k)$ is the Fourier transform of the convergence

κ_M provided by a subhalo of mass M,

$$\kappa_M(k) = \int \kappa_M(r) e^{ik\cdot r} d^2r = 2\pi \int \kappa_M(r) J_0(kr) r\, dr, \tag{2}$$

where the second equality holds for circularly symmetric $\kappa_M(r)$. Here, we make the flat-sky approximation, which is quite accurate given the $\sim$arcsecond field of view relevant for strong lensing.

Equation (1) is instructive in understanding exactly what aspects of the subhalo distribution control the form of the power spectrum shown in Figure 1. For example, note that on large scales (small wavenumber k), the substructure power spectrum plateaus to a constant value. The length scale above which $P(k)$ becomes flat corresponds to the sizes of the largest subhalos (compare blue vs. purple curves in the Figure). The amplitude of the power on these large scales is determined by the total abundance of subhalos of all masses, with a larger contribution from the most massive subhalos. This can be understood by inspecting Eqn. (1). Since $\kappa_M \propto M$, and assuming a power-law mass function $dn/dM \propto M^{-\alpha}$, then the integrand in Eqn. (1) behaves as $M^{3-\alpha}$, which is dominated by high masses for typical $\alpha \approx 2$. Towards smaller length scales, the power spectrum changes shape, declining towards higher k. The shape of the power spectrum on these scales is affected by two different terms: the internal profiles of massive halos, and the slope of the subhalo mass function (through the connection of tidal radius to subhalo mass). Fig. 1 illustrates the effects of varying either of these properties. Given a finite observable dynamic range, it may be difficult to disentangle these two effects.

3. THE LIKELIHOOD OF THE DENSITY POWER SPECTRUM

In this section, we describe the formalism for measuring the substructure power spectrum from lensing measurements. Suppose we have observations $\boldsymbol{O}$ (e.g., surface brightness maps) and random measurement noise $\boldsymbol{N}$ measured at n pixels. At each pixel, there is also a random deflection angle $\boldsymbol{\alpha}$ coming from substructure. We try to model the observations with a model that has parameters $\boldsymbol{p}$ describing the structure of the smooth lens potential and the background source emission. Suppose that both the noise $\boldsymbol{N}$ and deflections $\boldsymbol{\alpha}$ are Gaussian random fields with probability:

$$P(\boldsymbol{N}) = \frac{\exp\left(-\frac{1}{2}\boldsymbol{N}\cdot\mathbf{C}_N^{-1}\cdot\boldsymbol{N}\right)}{(2\pi)^{n/2}|\mathbf{C}_N|^{1/2}} \tag{3}$$

where $\mathbf{C}_N = \langle \boldsymbol{N}\,\boldsymbol{N}\rangle$ is the $n\times n$ noise covariance matrix, and similarly,

$$P(\boldsymbol{\alpha}) = \frac{\exp\left(-\frac{1}{2}\boldsymbol{\alpha}\cdot\mathbf{C}_\alpha^{-1}\cdot\boldsymbol{\alpha}\right)}{(2\pi)^{n}|\mathbf{C}_\alpha|^{1/2}} \tag{4}$$

where $\mathbf{C}_\alpha = \langle \boldsymbol{\alpha}\,\boldsymbol{\alpha}\rangle$ is the $2n\times 2n$ covariance matrix for deflection angles. Explicitly,

$$\langle \alpha_i(\boldsymbol{x})\,\alpha_j(\boldsymbol{x}+\boldsymbol{r})\rangle = A_1(r)\delta_{ij} + A_2(r)\frac{r_i r_j}{r^2} \tag{5}$$

$$A_1(r) = 4\int |\kappa(k)|^2 \frac{J_1(kr)}{kr}\frac{dk}{k}$$

$$A_2(r) = -4\int |\kappa(k)|^2 J_2(kr)\frac{dk}{k}.$$

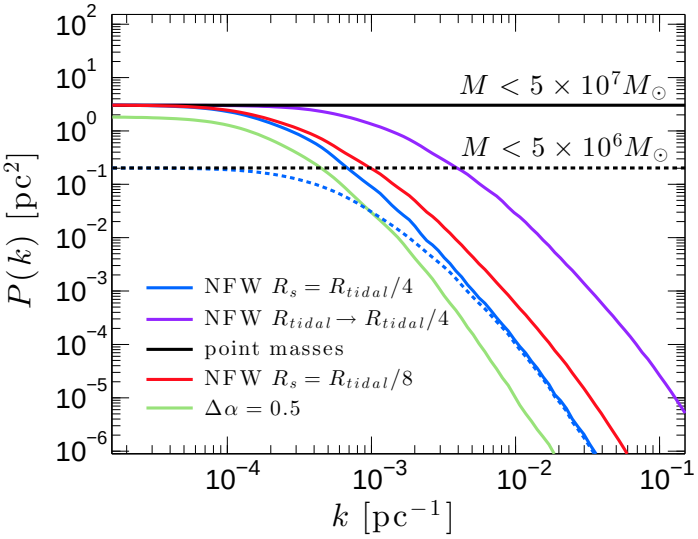

FIG. 1.— Power spectrum of projected density fluctuations from subhalos in the Via Lactea II (VL2) simulation. Subhalo masses, sizes, and locations in the VL2 catalog are used to generate theoretical power spectra using Eqn. (1). The blue solid curve shows our fiducial model which includes subhalos with $M < 5\times10^7 M_\odot$ with NFW profiles with $R_s = R_{tidal}/4$. The purple and red curves show the power spectrum when we alter the tidal radius, or the density profile (R_s) respectively. The solid black curve shows the power spectrum if the subhalos consist of point masses. The light-green curve shows the power spectrum when the slope of the mass function is altered by 0.5. The dotted lines show the power spectrum of subhalos with $M < 5\times10^6 M_\odot$, for our fiducial model (blue), and for the point mass model.

where we have used $\boldsymbol{\nabla}\cdot\boldsymbol{\alpha} = 2\kappa$. To estimate the likelihood for a given covariance given a set of measurements, we'll use Bayes' Theorem, which says that the likelihood for $\mathbf{C}_\alpha$, $\mathbf{C}_N$ is proportional to the likelihood for generating our observed measurements $\boldsymbol{O}_{\rm obs}$ given $\mathbf{C}_\alpha$ and $\mathbf{C}_N$:

$$\mathcal{L}(\boldsymbol{O}_{\rm obs},\boldsymbol{p}) = \int d^nN d^{2n}\alpha\, P(\boldsymbol{N})P(\boldsymbol{\alpha})$$
$$\delta\left[\boldsymbol{O}_m(\boldsymbol{p}) + \frac{\partial O}{\partial\alpha}\Delta\boldsymbol{\alpha} + \boldsymbol{N} - \boldsymbol{O}_{\rm obs}\right] P_{\rm p}(\boldsymbol{p}) \tag{6}$$

Here, $\boldsymbol{O}_m(\boldsymbol{p})$ is the model prediction for parameter set $\boldsymbol{p}$. Recall that $\boldsymbol{p}$ includes parameters for both the smooth lens and the source emission. In this work, we describe the source emission non-parametrically, as a pixelated map. Because the source map has many degrees of freedom that are not fully constrained by the observations, regularization is required to avoid over fitting (see e.g., Warren & Dye 2003; Suyu et al. 2006). This regularization acts as a prior, $P_{\rm p}(\boldsymbol{p})$, which multiplies the above likelihood. We use a Gaussian prior described by a covariance matrix $\mathbf{C}_{\rm p}$,

$$P_{\rm p}(\boldsymbol{p}) = \frac{\exp\left(-\frac{1}{2}(\boldsymbol{p}-\boldsymbol{p}_{\rm prior})\cdot\mathbf{C}_{\rm p}^{-1}\cdot(\boldsymbol{p}-\boldsymbol{p}_{\rm prior})\right)}{(2\pi)^{n_p/2}|\mathbf{C}_{\rm p}|^{1/2}}, \tag{7}$$

where n_p is the number of parameters, and $\boldsymbol{p}_{\rm prior}$ are fiducial parameters preferred by the prior. Without loss of generality, we will set $\boldsymbol{p}_{\rm prior} = 0$ to avoid confusion in the expressions below.

Assuming that the noise and substructure deflections are small, then the best-fitting parameters $\boldsymbol{p}$ are always

close to some fiducial parameter set $\boldsymbol{p}_0$. Taylor expanding, we have

$$\boldsymbol{O}_m(\boldsymbol{p}_0 + \Delta\boldsymbol{p}) \approx \boldsymbol{O}_0 + \frac{\partial O}{\partial p}\Delta\boldsymbol{p} + \ldots, \tag{8}$$

By marginalizing over the uncertain model parameters (smooth lens and source parameters) we can calculate the marginalized likelihood of the substructure covariance matrix. After a few lines of algebra we arrive at:

$$\mathcal{L}(C_\alpha) = (|C_N|\,|C_\alpha|\,|C_p||M|)^{-1/2} e^{\frac{1}{2}B^\mathsf{T}\,M\,B} \\ e^{-\frac{1}{2}(\Delta\boldsymbol{O}^\mathsf{T} C_N^{-1}\Delta\boldsymbol{O} + \boldsymbol{p}_0 C_p^{-1}\boldsymbol{p}_0)} \tag{9}$$

where

$$M = \begin{bmatrix} C_\alpha^{-1} + \frac{\partial O}{\partial\alpha}^\mathsf{T} C_N^{-1}\frac{\partial O}{\partial\alpha} & \frac{\partial O}{\partial\alpha}^\mathsf{T} C_N^{-1}\frac{\partial O}{\partial p} \\ \frac{\partial O}{\partial p}^\mathsf{T} C_N^{-1}\frac{\partial O}{\partial\alpha} & C_p^{-1} + \frac{\partial O}{\partial p}^\mathsf{T} C_N^{-1}\frac{\partial O}{\partial p} \end{bmatrix} \tag{10}$$

and

$$B = \begin{bmatrix} \Delta\boldsymbol{O}^\mathsf{T} C_N^{-1}\frac{\partial O}{\partial\alpha} & \Delta\boldsymbol{O}^\mathsf{T} C_N^{-1}\frac{\partial O}{\partial p} - \boldsymbol{p}_0^\mathsf{T} C_P^{-1} \end{bmatrix} \tag{11}$$

Here C_p is the prior covariance matrix (regularization matrix). By Bayes' theorem, $\mathcal{L}(\mathbf{C}_\alpha, \mathbf{C}_N) \propto \mathcal{L}(\boldsymbol{O}_{\rm obs})$. Thus, by mapping out $\mathcal{L}(\boldsymbol{O}_{\rm obs})$ as a function of $\mathbf{C}_\alpha$ and $\mathbf{C}_N$, we can determine the likelihood of the noise level and substructure power spectrum for each set of observations.

4. SIMULATIONS

We generate mock observations of galaxies lensed by a halo (macro lens) and a population of subhalos, and use Equation (9) to map the likelihood of the amplitude of the power spectrum of the subhalo field, $P(k)$. The macro lens is modeled as a power-law elliptical mass distribution (Barkana 1998) plus two additional angular multipoles ($\cos 3\theta$ and $\cos 4\theta$), along with external shear. The substructure population is modeled as a Gaussian random density field with a given power spectrum. To calculate the deflections due to the subhalo field, a map of substructure surface density is generated and the deflection angles are calculated in Fourier space using:

$$\tilde{\alpha} = \left(\frac{2\,i\,k_x}{k^2}\tilde{\kappa}, \frac{2\,i\,k_y}{k^2}\tilde{\kappa}\right) \tag{12}$$

where $\tilde{\kappa}$ is the Fourier transform of the density field, and k_x and k_y are the Fourier coordinates. The surface density, κ, is generated by an inverse Fourier transformation of a map whose real and imaginary components are drawn at random to give rise to a desired power spectrum. To avoid periodicity and edge effects, we construct a subhalo density map that is approximately ten times larger than the $4'' \times 4''$ field of interest. The deflection angles due to the main lens and the substructure field are then added together to predict the lensed images of a background source. The source consists of multiple (1-5) star-forming clumps with a Gaussian light profile with FWHM of ~ 300 pc, distributed over an area of $\sim 1-2$ kpc. As discussed in the previous section, we have used a pixelated grid to parameterize the background source emission, whose reconstruction is regularized using a Gaussian prior. The source pixels covered an area of 3×3 kpc in the source plane with 40×40 pixels. Finally, Gaussian noise was then added to the lensed images, at levels based on the signal to noise of previous ALMA observations of lensed dusty sources.

For high-excitation molecular transitions it is expected that the source emission will be composed of a number of discrete clumps embedded in a larger structure, such as an exponential disk. We used this structure to construct the source prior by calculating the power spectrum and covariance of such a clustered source model. The power spectrum of the source emission can be calculated in the same way as our halo model approach for computing the lens substructure power spectrum. Suppose that we have N_c clumps in our source galaxy, whose distribution within the galaxy has profile $U_c(r)$. We normalize U_c to have unit integral, $\int U_c(r) d^2r = 1$. Its Fourier transform is $U_c(k)$. Clump i has luminosity L_i and profile $u_i(r)$, normalized to have unit integral, $\int u_i(r) d^2r = 1$. Then the power spectrum of the source emission is proportional to

$$P_{\rm src}(k) \propto \left[\sum_i^{N_c} L_i^2 |u_i(k)|^2 + \sum_{i\neq j}^{Nc} L_i L_j |U_c(k)|^2 u_i^*(k) u_j(k)\right] \tag{13}$$

To find the overall normalization of the covariance matrix we used the method presented in Suyu et al. (2006) to maximize the evidence. Note that this clump model is only used to construct the covariance matrix that regularizes the source reconstruction. We do not assume that the source is clumpy, but instead allow an arbitrary source emission. This clump model is only used to regularize the pixelated source reconstruction. In agreement with previous work, we have found that for data with sufficiently high signal-to-noise, the reconstruction does not depend sensitively on the precise form of the regularization.

To map the likelihood of the power spectrum amplitude, we construct the deflection covariance matrix, $\mathbf{C}_\alpha$, using equation (5). We use finite differencing to construct $\partial O/\partial p$ and $\partial O/\partial\alpha$, and use the macro model (without including the substructure field) as the reference model, O_0.

The size of these matrices and the computational cost of inversions grow very rapidly with increasing image resolution. For an image with $n \times n$ pixels, the deflection covariance matrix has a size of $2n^2 \times 2n^2$. Inversion of this matrix has a typical time complexity of $\sim \mathcal{O}((2n^2)^{2.5})$. In other words, doubling image resolution results in a $\sim$ 32-fold increase in computational costs. A single evaluation of the likelihood (which includes multiple inversions and determinant calculations) for a 50×50 image on a single CPU could take up to a few minutes. At higher resolutions, not only the likelihood evaluation becomes remarkably slower, retaining the data on single machine memory becomes unfeasible, with the overall size of matrices exceeding tens of GB for a 100×100 image. To overcome these obstacles we use the *Elemental* package, an open-source `C++` library for distributed-memory dense linear algebra (Poulson et al. 2013; Petschow et al. 2012).

In the simulations presented in this work we have assumed CCD data with uncorrelated noise in the

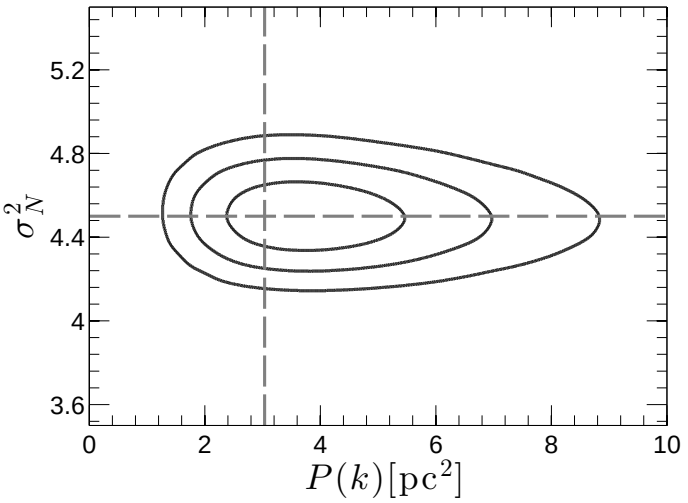

FIG. 2.— Joint-likelihood of noise and the amplitude of the power spectrum, mapped by evaluating Equation (9) using mock observations described in §4, lensed by a density field which includes substructure with a flat power spectrum. The dashed lines show the true values which were used in the mock observation. The input amplitude of the power spectrum is successfully recovered, with little if any degeneracy between instrumental noise and substructure fluctuations.

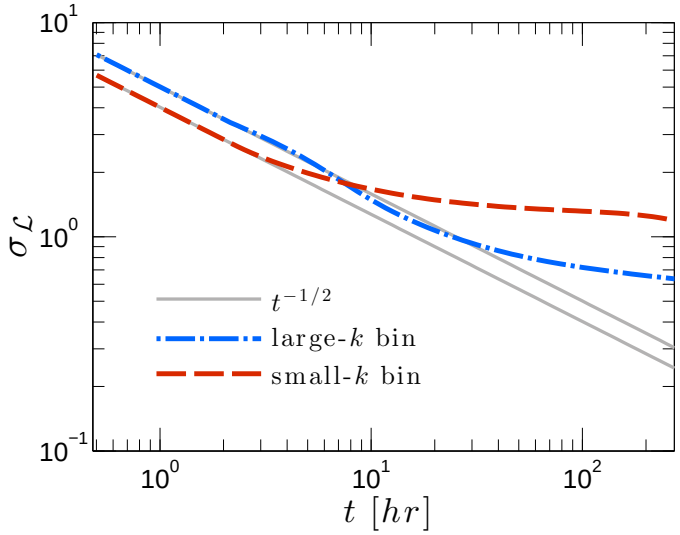

FIG. 3.— The rms of the likelihood of the amplitude of the power spectrum as a function of observing time. For purely noise dominated observations, these curves should scale as $t^{-1/2}$, as seen for observations $\lesssim 10$ hours. For longer observations, the errors are dominated by sample variance. The contribution of sample variance is larger for modes with smaller wavenumber k, (red) compared to the more numerous modes at higher frequency (blue).

images. This method, however, is equally applicable to interferometric data using the same fitting procedure as in Hezaveh et al. (2013a) and Hezaveh et al. (2013b). In that case, the observables are the measured visibilities and, in the uv-space, the noise covariance matrix $\mathbf{C}_N$ is diagonal.

5. ANALYSIS AND DISCUSSIONS

5.1. *Substructure vs. other sources of fluctuations*

Density perturbations produce random fluctuations in the lensed observables. Instrumental noise also produces random errors in the observables, and naively we might expect that the effects of measurement noise could be difficult to disentangle from the effects of mass substructure. Similarly, we might worry that fluctuations in the source emission might also be degenerate with substructure fluctuations in the lens mass distribution. Figure 2 shows that this is not the case. The figure plots the joint likelihood (from Equation 9) as measured from simulated observations of sources lensed by a main halo and a density field with a flat, white noise power spectrum. As described above, we have marginalized over a pixelated source emission grid. As is apparent, there is little if any degeneracy between instrumental noise and the amplitude of the density power spectrum. The density field in this simulation only contains modes between $\sim 0.04 - 0.4$ kpc^{-1}, which roughly cover the range of modes where the power spectrum is expected to be flat (see Figure 1). Repeating this procedure for various mock observations indicates that the two parameters are non-degenerate over the entire range that we have simulated. We again stress that this calculation marginalizes over the source emission, so the fact that we recover the input power spectrum implies that substructure in the lens mass is not degenerate with clumpiness in the source emission. We have evaluated this likelihood using different source priors with different parameters (i.e. U_c and u in Equation 13) and found that for high signal to noise observations the precise form of the prior does not appear to significantly affect the reconstruction.

Given the lack of any degeneracy between measurement noise and lens substructure, and the fact that the noise properties of most observations can be precisely quantified, in the rest of this work we do not map the likelihood along the noise dimension, assuming that the noise rms is accurately known.

5.2. *Sample variance*

Next, we investigate how the errors of the power spectrum measurement scale with signal to noise. Our results suggest that the power spectrum uncertainty becomes sample variance dominated for observations with very high signal to noise. The red and blue dashed curves in Figure 3 show the rms of likelihood as a function of observing time for continuous observations. The blue dashed curve corresponds to the rms for modes with larger wavenumber than the ones for the red curve. As seen in the figure, at higher wavenumber (blue curve), each bin contains a larger sample of modes, helping to reduce sample variance. The gray lines show the $t^{-1/2}$ scaling expected for measurements that are purely noise dominated. To overcome the limitation imposed by sample variance when higher precision measurements are needed, we can combine measurements from many lenses, assuming that they are different realizations of the same process. This approach, however, requires a careful analysis of the selection methods and the scaling of the subhalo properties with that of the host halo (Xu et al. 2013).

5.3. *Measuring the shape of $P(k)$*

So far, we have discussed measurements of flat (white noise) power spectra. This power spectrum arises from populations of point masses (e.g. primordial black holes).

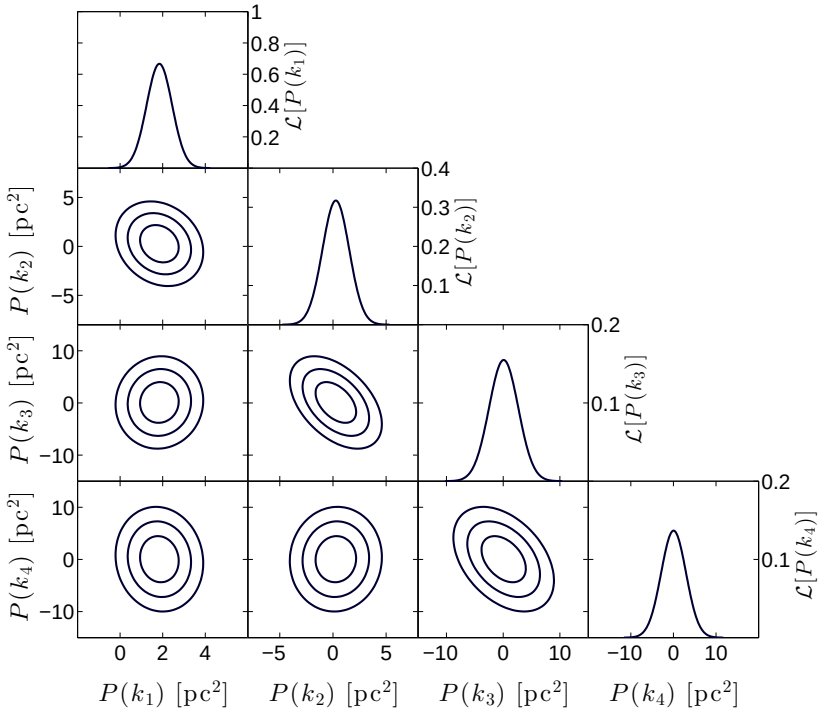

FIG. 4.— Fisher matrix forecasts for $P(k)$ errors for a 10 hour long observation. Here we have parametrized the power spectrum with bandpowers in 4 bins, k_1 to k_4, corresponding to modes in the ranges $0.04-0.4$, $0.4-0.8$, $0.8-1.2$, and $1.2-1.7$ kpc^{-1} respectively. Negative allowed values of $P(k)$ are, of course, an artifact of the Fisher matrix approximation to the full likelihood.

As discussed in section 2, realistic dark matter subhalos, with finite sizes and smooth density profiles, give rise to power spectra that are flat on large scales (small k) and fall off at large wavenumbers (Figure 1). To measure a power spectrum with an arbitrary shape, we can parameterize $P(k)$ in Equation 5 as a function of k when constructing $\mathbf{C}_\alpha$, and fit for the free parameters using Equation 9. For example, we could assume that the power spectrum is piecewise flat in discrete k bins, and then use the bandpowers in those bins as the free parameters. For sufficiently narrow bins, we can approximate arbitrary power spectra this way. Noting that the power is relatively flat over $0.04-0.4\,\mathrm{kpc}^{-1}$, we choose a linear binning in which the first bin covers the range of k modes over the flat part of the power spectrum, and the other bins measure the fall of the power spectrum at high k. With this choice of binning, a measurement of the amplitude of the modes in the first bin gives power over the flat part of the power spectrum, revealing the total abundance of all low-mass subhalos.

For n bins over the available range of k's, the power spectrum is defined in an n-dimensional parameter space. Since the likelihood evaluations are computationally expensive, we use a Fisher analysis to forecast the size of the errorbars and the degeneracies between power at different scales for different observing conditions. The input power spectrum of the subhalo density field in the mock observations is set to be consistent with the Via Lactea II (VL2) simulation: the positions, masses, and tidal radii of the subhalo are taken from the publicly available VL2 catalogue[9] (Diemand et al. 2008) and the subhalos are given a truncated NFW profile with $R_s = R_{tidal}/4$. Figure 4 shows an example of the parameter covariance (amplitudes in four bins) for a simulated observation.

Figure 5 shows the errorbars of two bins for a signal to noise comparable to a 10-hr long ALMA observation of bright lensed dusty galaxies. This results in a detection of the power in the first bin ($\sim 3\sigma$) revealing the total abundance of subhalos. On smaller scales, the predicted power spectrum falls too rapidly and this observation can only put an upper bound on the high-k amplitude. This upper limit, however, may be adequate to indicate a break in the power spectrum. An observation approximately 4 times longer (or involving 4 different lens systems) could measure the power over this regime.

5.4. *Non-Gaussianity*

So far, we have assumed that the subhalo density field could be treated as a Gaussian random field. In reality, the substructure field is not Gaussian distributed. The non-Gaussianity mainly arises from the few most massive subhalos. To reduce this non-Gaussianity, it is important to be able to detect and remove the effect of the most massive subhalos with low number densities. The power spectra used for simulations in this work were calculated for subhalos with $M < 5\times10^7 M_\odot$, assuming that subhalos with masses larger than this limit could be detected individually using a direct lens modeling approach (e.g. Vegetti et al. 2012; Hezaveh et al. 2013a). To estimate how much the remaining non-Gaussianity in the density field biases our results, we performed 100 simulations of Gaussian and non-Gaussian substructure density fields. The non-Gaussian maps were generated with subhalo masses and numbers taken from the Via Lactea II catalogue. After mapping the likelihood of the power spectrum amplitude for each simulation,

[9] http://www.ucolick.org/∼diemand/vl/data.html

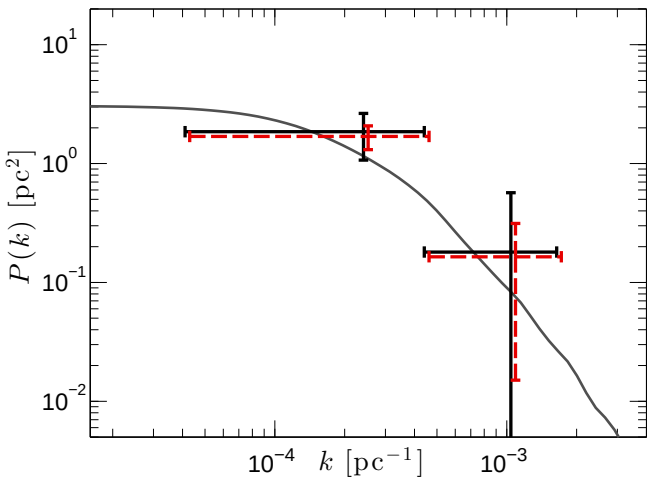

FIG. 5.— Forecast for a measurement of the power spectrum of subhalos with $M < 5{\times}10^7\,M_\odot$ for a 10-hour long observation (black errorbars) of a single source with ALMA, assuming an observed source continuum flux of 50 mJy at 850 μm. The first bin, with a significance of $\sim 3\sigma$, indicates the abundance of all subhalos in the main dark matter halo. Deeper observations ($\sim$ 40 hr), combination of all the modes at higher k, and more favorable conditions (smaller source size) could allow a measurement of the break in the power spectrum at higher k (red errorbars). The underlying power spectrum is the fiducial model of Figure 1 (blue curve in Figure 1).

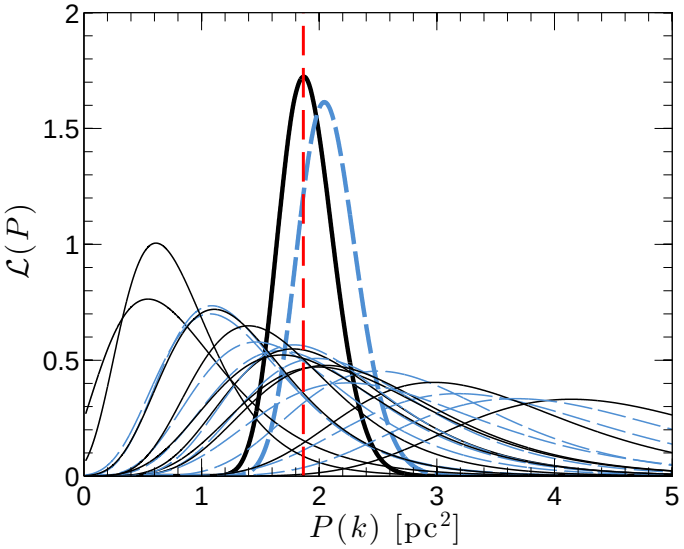

FIG. 6.— The thin curves show the recovered likelihoods of the amplitude of the power spectrum in ten realizations of a Gaussian (solid black) and non-Gaussian (dashed blue) density field. In the case of blue curves, the subhalo density field is non-Gaussian, but incorrectly treated as a Gaussian field. The thick black solid curve show the combined likelihood that results from multiplying the ten Gaussian likelihoods together. The thick blue dashed curve shows the same for the non-Gaussian density maps. The deviation of the blue dashed curve from the solid curve ($\sim$ 7%) results from biasing the likelihood estimator due to non-Gaussianity in the substructure field. Repeating the test for 100 realizations yields consistent results, with a corresponding improvement in the joint likelihood.

we multiplied the hundred likelihoods together for the Gaussian and non-Gaussian case. Figure 6 shows a subset of the resulting likelihoods. The modes depicted in the plot correspond to the first bin of Figure 5, covering the flat part of the spectrum. This figure shows that in the case of non-Gaussian density fields, the true value is about 7% biased. Although this is a biased measurement, given that we currently do not know the value of this power spectrum to any precision, a 7% biased measurement is valuable. However, it is also possible to avoid this bias, in principle.

If we can assume a known profile for subhalos (e.g. truncated NFW), then we can generate Monte Carlo realizations of the non-Gaussian density field and constrain substructure properties using the method of Dalal & Kochanek (2002). Additionally one can perform such analysis for different density profiles and marginalize over their parameters. That method, however, is considerably more computationally intensive than the Gaussian likelihood estimator discussed in this paper.

6. CONCLUSIONS

In this paper, we presented a framework for measuring the power spectrum of the substructure density field using observations of strong gravitational lenses. We showed how the amplitude and shape of the power spectrum is related to the abundance, density profile, and mass function of subhalos. Using mock observations, we tested the method, successfully recovering the input parameters and showed that if dark matter halos host large populations of subhalos consistent with CDM simulations, this power spectrum could be measured with near-future observations.

We found that $\sim$ 10 hour long ALMA observations of a single lensed submm source should be capable of detecting the amplitude of the substructure power spectrum at $\gtrsim 3\sigma$ significance. In our calculations, we have marginalized over any uncertainty in the 'macro' model describing the smooth mass distribution of the main lens. These macro parameters are completely degenerate with the longest wavelength modes of the substructure field, implying that the lowest k modes will be unconstrained. Fortunately, however, such unconstrained modes are few in number, and their degeneracy does not significantly impact on our ability to measure low k bandpowers, as long as sufficiently wide k-bins are employed, as illustrated in Figs. 4 and 5.

The power spectrum measured from lensing observations may be directly compared to results of numerical simulations. Although we have interpreted the power spectrum in terms of the abundance of dark matter subhalos, the same quantity measured by lensing may also directly be measured in simulations without resorting to catalogs of subhalos. This may prove to be a useful approach, since subhalo properties are notoriously difficult to measure in simulations: different subhalo finders applied to the same simulations can sometimes produce discrepant results, depending on the subhalo definitions and parameter choices adopted by the various finders (Onions et al. 2012). By directly measuring substructure power spectra from simulations, uncertainties in the definition of subhalos may be circumvented.

Lastly, we note that our results depend on our choice of parameters in our simulations. Wherever possible, we have attempted to be conservative in our

choices. We have generated macro lenses and source brightnesses consistent with existing low-resolution imaging of sub-mm lenses from ALMA (Hezaveh et al. 2013b; Vieira et al. 2013). The main uncertainty in our simulations is the unknown number and size distribution of star-forming clumps in the source galaxies. To be conservative, we have assumed clump sizes of $\sim$ 200 pc, the upper limit placed by current observations (e.g. Swinbank et al. 2010). If the source clumps are smaller in reality than we have assumed, then our forecasted constraints on the high-k power spectrum can improve significantly.

This work was supported by NASA under grant NNX12AD02G. YH thanks the *Elemental* team, Jack Poulson and Jeff Hammond for tremendous support. YH acknowledges very useful and constructive discussions with Ryan Keisler, Phil Marshall, Roger Blandford, James Bullock, and David Spergel. This work was initiated at the Aspen Center for Physics, which is supported by NSF grant 1066293. We thank Calcul Quebec for providing the computing resources that we used for a part of this research.

REFERENCES

Barkana, R. 1998, ApJ, 502, 531, astro-ph/9802002
Bussmann, R. S., Pérez-Fournon, I., Amber, S., et al. 2013, ApJ, 779, 25, 1309.0836
Chamberlain, R. T., Dalal, N., Hearin, A., & Ricker, P. 2014, ArXiv e-prints, 1407.2648
Cooray, A., & Sheth, R. 2002, Phys. Rep., 372, 1, astro-ph/0206508
Dalal, N., & Kochanek, C. S. 2002, ApJ, 572, 25, astro-ph/0111456
Diemand, J., Kuhlen, M., & Madau, P. 2007a, ApJ, 657, 262, astro-ph/0611370
—. 2007b, ApJ, 667, 859, arXiv:astro-ph/0703337
Diemand, J., Kuhlen, M., Madau, P., et al. 2008, Nature, 454, 735, 0805.1244
Diemand, J., Moore, B., & Stadel, J. 2004, MNRAS, 352, 535, astro-ph/0402160
Hezaveh, Y., Dalal, N., Holder, G., et al. 2013a, ApJ, 767, 9, 1210.4562
Hezaveh, Y. D., Marrone, D. P., Fassnacht, C. D., et al. 2013b, ApJ, 767, 132, 1303.2722
Klypin, A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, ApJ, 522, 82, astro-ph/9901240
Kravtsov, A. 2010, Advances in Astronomy, 2010, 8, arXiv:0906.3295
Kravtsov, A. V., Gnedin, O. Y., & Klypin, A. A. 2004, ApJ, 609, 482, astro-ph/0401088
Lovell, M. R., Frenk, C. S., Eke, V. R., et al. 2014, MNRAS, 439, 300, 1308.1399
Lovell, M. R., Eke, V., Frenk, C. S., et al. 2012, MNRAS, 420, 2318, 1104.2929
Mao, S., & Schneider, P. 1998, MNRAS, 295, 587, arXiv:astro-ph/9707187
Moore, B., Ghigna, S., Governato, F., et al. 1999, ApJ, 524, L19, astro-ph/9907411
Navarro, J. F., Ludlow, A., Springel, V., et al. 2010, MNRAS, 402, 21, arXiv:0810.1522
Nierenberg, A. M., Treu, T., Wright, S. A., Fassnacht, C. D., & Auger, M. W. 2014, ArXiv e-prints, 1402.1496
Onions, J., Knebe, A., Pearce, F. R., et al. 2012, MNRAS, 423, 1200, 1203.3695
Petschow, M., Peise, E., & Bientinesi, P. 2012, CoRR, abs/1205.2107
Poulson, J., Marker, B., van de Geijn, R. A., Hammond, J. R., & Romero, N. A. 2013, ACM Trans. Math. Softw., 39, 13:1
Rocha, M., Peter, A. H. G., Bullock, J. S., et al. 2013, MNRAS, 430, 81, 1208.3025
Simon, J. D., & Geha, M. 2007, ApJ, 670, 313, 0706.0516
Stadel, J., Potter, D., Moore, B., et al. 2009, MNRAS, 398, L21, 0808.2981
Strigari, L. E., Bullock, J. S., Kaplinghat, M., et al. 2007, ApJ, 669, 676, 0704.1817
Suyu, S. H., Marshall, P. J., Hobson, M. P., & Blandford, R. D. 2006, MNRAS, 371, 983, astro-ph/0601493
Swinbank, A. M., Smail, I., Longmore, S., et al. 2010, Nature, 464, 733, 1003.3674
Vegetti, S., Koopmans, L. V. E., Bolton, A., Treu, T., & Gavazzi, R. 2010, MNRAS, 408, 1969, 0910.0760
Vegetti, S., Lagattuta, D. J., McKean, J. P., et al. 2012, Nature, 481, 341, 1201.3643
Vieira, J. D., Marrone, D. P., Chapman, S. C., et al. 2013, Nature, 495, 344, 1303.2723
Viel, M., Weller, J., & Haehnelt, M. G. 2004, MNRAS, 355, L23, astro-ph/0407294
Warren, S. J., & Dye, S. 2003, ApJ, 590, 673, astro-ph/0302587
Xu, D. D., Sluse, D., Gao, L., et al. 2013, ArXiv e-prints, 1307.4220